

Reconfigurable Low Energy Multiplier for Multimedia System Design

Suhwan Kim and Marios C. Papaefthymiou
Advanced Computer Architecture Laboratory
Department of Electrical Engineering and Computer Science
University of Michigan
Ann Arbor, MI 48109
suhwan.marios@eecs.umich.edu

Abstract

This paper proposes a reconfigurable pipelined multiplier architecture that achieves high performance and very low energy dissipation by adapting its structure to computational requirements over time. In this reconfigurable multiplier, energy is saved by disabling and bypassing an appropriate number of pipeline stages whenever input data rates are low. To evaluate the efficiency of our multiplier architecture, we have designed a multiplier-based inverse quantizer (IQ) for MPEG-2 MP@ML. Pipelines are dynamically reconfigured according to the size of the picture and the number of nonzero quantized DCT coefficients per block. In comparison with corresponding multiplier implementations that use conventional pipelines, our reconfigurable multipliers dissipate about 31–58% less energy. Relative energy savings increase with decreasing data rates, since our reconfigurable structures stay in a low energy configuration for proportionately longer time.

1 Introduction

Multimedia wireless communications have resulted in a growing demand for energy-efficient video processing. Next generation portable devices must provide support for low-energy encoding/decoding and transmission of multimedia information. Several video standards are currently in use, including MPEG-1 and MPEG-2 for multimedia applications, and H.261 for video-phone and video-conferencing applications. Implementations of these standards for mobile system-on-chip devices should provide substantial computing capabilities at low energy consumption levels [1, 8, 9].

Multiplication is a key arithmetic operation in video processing. The development of multipliers with short critical paths and low power consumption has become the topic of extensive recent investigation [2, 5, 6, 12]. Pipelining is

a popular technique for realizing high-performance, high-efficiency CMOS multipliers by reducing the supply voltage at the lowest possible level while still satisfying throughput constraints. In deep pipelines, however, registers are responsible for an increasingly large fraction of total dissipation, no matter how efficiently they may have been implemented [4, 7, 10, 13]. Even if clock gating is used to only store essential data, pipeline registers may still be latching their inputs unnecessarily if throughput requirements are lower than the maximum specified.

This paper presents a reconfigurable pipelined multiplier that adapts its performance and dissipation to its computational load over time. Our multiplier is capable of adapting its structure within one clock cycle, if required. It can thus efficiently cope with variable data-rate multimedia applications such as video processing. Energy dissipation is reduced by disabling and bypassing a select subset of registers based on the specified throughput requirements and the anticipated computational load. This information is application-dependent and can be inferred at high abstraction levels. In the context of inverse quantization for video processing, for example, our multiplier can be adapted by simply counting the number of nonzero coefficients in each encoded block. Our reconfiguration approach can be applied to general linear pipelines. It also can be combined with voltage scaling to further increase energy efficiency.

To evaluate the efficiency of our reconfigurable multiplier architecture, we designed a multiplier-based inverse quantizer (IQ) for MPEG-2 MP@ML. Pipelines were dynamically reconfigured according to the picture size and the number of nonzero quantized discrete cosine transform (DCT) coefficients per block. In simulations with a 0.35 μm standard-cell CMOS technology and MPEG-2 MP@ML bitstreams, our reconfigurable IQ was up to 58% less dissipative than its non-reconfigurable, statically pipelined counterpart. Moreover, relative reductions increased as data rates decreased, since the reconfigurable multiplier stayed in a low energy mode for proportionally longer time.

The remainder of this paper has six sections. In Section 2 we briefly review the MPEG video processing standard. Section 3 gives background on multiplier-based IQ design and motivates the use of reconfigurable multipliers for improving the energy efficiency of IQ's. Section 4 describes our design methodology for high-performance, low-energy reconfigurable multipliers. Section 5 presents our comparative evaluation of IQ's for a variety of bitstreams and bit rates. Our contributions and ongoing research are summarized in Section 6.

2 MPEG Video Processing

In the MPEG video standard, a coded frame consists of a frame picture or a pair of field pictures. There are three types of pictures: intra (I), predictive (P), and bi-directional (B). Each picture is divided into non-overlapping macroblocks with 16×16 pixels. Each macroblock (MB) consists of four luminance blocks (Y) and two chrominance blocks (Cb and Cr), each 8×8 pixels. Since I-pictures are coded without reference to neighboring pictures in the sequence, their coding exploits only the correlations within the picture. P-pictures and B-pictures are coded as differences between the picture being coded and a reference picture. If there is motion in the sequence, a better prediction can be obtained from pixels in the reference picture that are shifted relative to the current picture pixels.

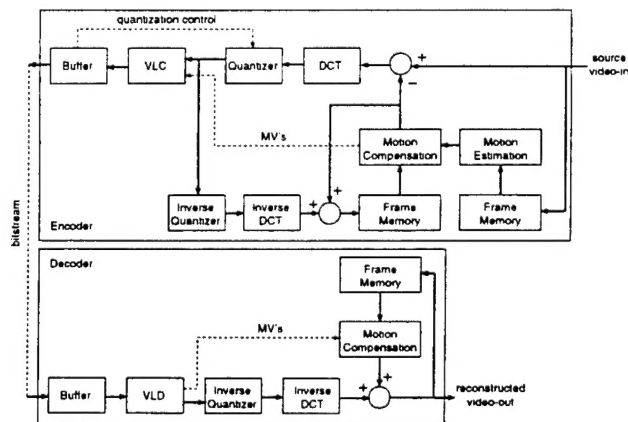


Figure 1. Block diagram of basic MPEG video encoder and decoder.

The block diagram of the basic MPEG video encoder and decoder structure is shown in Figure 1. In general, the encoder comprises a DCT/IDCT, a motion estimator/compensator, a Q/IQ, and a variable length coder (VLC). The decoder performs the reverse operations of the encoder and consists of a variable-length decoder (VLD), an IQ, an IDCT, and a motion compensator.

Table 1. Upper bounds for picture size, frame rates, and bit rates.

PROFILE LEVEL		Simple	Main	SNR	Spatial	High
Low	pels*2/frame frames/sec Mbit/sec		352 / 288 30 4	352 / 288 30 4		
Main	pels*2/frame frames/sec Mbit/sec	720 / 576 30 15	720 / 576 30 15	720 / 576 30 15		720 / 480 30 20
High-1440	pels*2/frame frames/sec Mbit/sec		1440 / 1152 60 60		1440 / 1152 60 60	1440 / 1152 60 80
High	pels*2/frame frames/sec Mbit/sec		1920 / 1152 60 60			1920 / 1152 60 100

MPEG video standards only specify the video bitstream syntax, the decoding semantics, and the required maximum performance for decoding a bitstream of any particular type. Table 1 gives the upper bounds specified by MPEG for picture size, frame rate, and bit rate for various combinations of profiles and levels. For example, for the main-profile/main-level combination (MPEG-2 MP@ML for short) in the NTSC-compatible mode, the picture size, frame rate, and bit rate are 720×480 pixels, 30 frames/sec, and 15 Mbit/sec, respectively. Thus, the maximum throughput requirement for an MPEG-2 MP@ML decoder is 15.552×10^6 samples/sec ($30 \text{ frames/sec} \times (720/16 \times 480/16) \text{ MB/frame} \times 6 \text{ blocks/MB} \times 64 \text{ pixels/block}$). Therefore, each building block in Figure 1, such as the IQ for example, should be designed to meet this maximum throughput requirement, even though the average data rate is usually much less than the upper bound specified by the standard. Consequently, the decoder may be performing no useful computation for a large fraction of the total number of cycles.

3 Inverse Quantization

In this section, we describe the inverse quantization procedure in the MPEG video standard, following largely the description given in [8]. We subsequently motivate the use of reconfigurable multipliers for reducing its energy dissipation.

The DCT-based coding/decoding for an 8×8 block of pixels is common to all picture types and plays a central role in the MPEG video standard. The DCT has certain properties that simplify coding models and make coding coefficients using perceptual quality measures. Basically, the DCT is a method for decomposing a block of pixels into a weighted sum of spatial frequencies. If only the low frequency DCT coefficients are nonzero, the pixels in the block vary slowly with position. If high frequencies are present, the block intensity changes rapidly from pixel to

pixel. When the DCT is computed for a block of pixels, it is desirable to represent the coefficients for high spatial frequencies with less precision. This is done by a process called quantization. A DCT coefficient is quantized by dividing it by a nonzero positive integer called a quantization value and rounding the quotient—the quantized DCT coefficient—to the nearest integer. The bigger the quantization value is, the lower precision coefficients can be transmitted to a decoder with fewer bits. The use of large quantization values for high spatial frequencies allows the encoder to selectively discard high spatial frequency activity that the human eye cannot readily perceive. The quantizing values are chosen so as to minimize perceived distortion in the reconstructed pictures, using principles based on the human visual system.

The coding of quantized DCT coefficients is lossless, that is, the decoder is able to reproduce the exact same DCT coefficients computed by the encoder, before 2-D IDCT is processed. This process is essentially a multiplication by the quantizer step size. For each block, the quantizer step size is determined by the product of a weighting matrix and a quantizer scale factor. The following equation specifies how to reconstruct the DCT coefficients from the quantized ones.

$$F(v, u) = (2 \cdot QF(v, u) + k) \cdot W(w, v, u) \cdot S_q / 32,$$

where $F(u, v)$ is an 8×8 matrix of reconstructed DCT coefficients, $QF(u, v)$ is the corresponding quantized DCT coefficient matrix, $W(w, v, u)$ is a weighting matrix ($w = 0$ for intra coded blocks, $w = 1$ for non-intra coded blocks), k is a parameter ($k = 0$ for intra coded blocks, $k = \text{Sign}(QF(v, u))$ for non-intra coded blocks), and S_q is the quantizer scale.

DCT blocks of MPEG compressed video sequences usually have only five or six nonzero coefficients, mainly located in the low spatial frequency position. Given such input data statistics, the number of operations per block can be reduced, since multiplication and addition with a zero-valued DCT coefficient constitute no operation. The number of operations required for each block can be predicted precisely and effortlessly, before the IQ is started, by observing the operation of the VLD. To achieve the maximum IQ throughput for MPEG-2 MP@ML (15.552×10^6 samples/sec), deeply pipelined multipliers can be used. When there are many zero coefficients, however, this peak performance is not required and results in wasted energy.

Reconfigurable pipelined multipliers can save energy by adapting to the throughput needs of the input stream over time. The number of nonzero quantized DCT coefficients per block equals the number of operations required to perform the block's inverse quantization. As the number of nonzero quantized DCT coefficients per block decreases, the throughput of the multiplier also decreases, thus saving

energy. Picture size can be used to achieve additional energy savings. For example, if the picture size is a quarter of the maximum picture size for MPEG-2 MP@ML, the number of blocks that must be processed within 1/30 second decreases by a factor of 4. Therefore, the multiplier throughput may decrease by a factor of 4, independent of the number of nonzero quantized DCT coefficients per block.

4 Reconfigurable Multiplier Design

This section highlights our methodology for designing performance-driven reconfigurable multipliers. In our proposed reconfigurable structure, whenever throughput requirements are low, register stages are selectively disabled by gated clocks and bypassed by multiplexors.

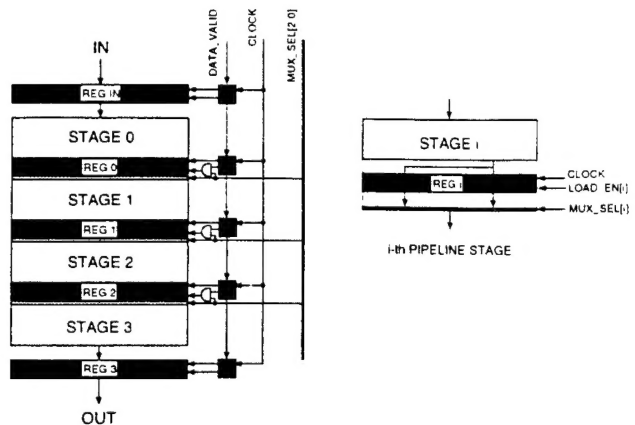


Figure 2. Reconfigurable 4-stage pipeline.

Figures 2 shows our 4-stage pipelined reconfigurable structure. The throughput of a conventional pipelined structure is fixed at one operation per cycle, whereas the throughput of our reconfigurable pipelined structure may be set to one operation every one, two, or four cycles, depending on the input data rates.

Figure 3 shows the timing diagram of a non-reconfigurable pipeline that processes 2 samples over 8 clock cycles. Figure 4 gives the timing diagram of the reconfigurable 4-stage pipeline, configured in its single-stage mode. The 4-stage pipeline is capable of handling the maximum throughput requirement of 8 samples per 8 cycles. When only 2 samples need to be processed in 8 cycles, however, the conventional 4-stage pipeline remains idle for 6 cycles. The reconfigurable pipeline of Figure 2 uses these idle cycles to spread the computation and eliminate three stages of registers. Three register stages are disabled by gated clocks and bypassed through multiplexors, thus saving a significant fraction of the datapath's total dissipation in the reconfigured datapath.

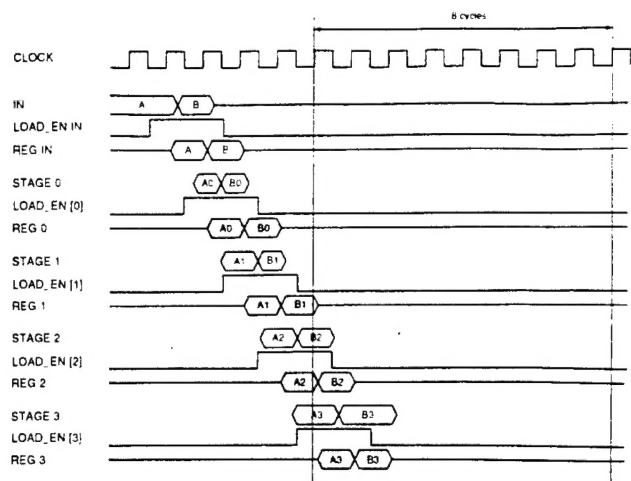


Figure 3. Timing diagram of non-reconfigurable 4-stage pipeline for 2 samples over 8 cycles.




Figure 4. Timing diagram of reconfigurable 4-stage pipeline in single-stage mode for 2 samples over 8 cycles.

Array multiplier structures are popular due to their simple and regular interconnections. Figure 5 shows a 4-stage pipelined 4×4 array multiplier which is designed using our reconfiguration technique. This multiplier can be reconfigured as a 1, 2, and 4-stage pipeline according to the throughput desired. We designed, synthesized and simulated 16×16 multipliers with different numbers of pipeline stages using a $0.35 \mu\text{m}$ CMOS standard-cell technology. The registers of the reconfigurable multipliers were implemented by positive edge-triggered D flip-flops using transmission gates. This kind of flip-flops is commonplace in standard cell design [11]. To obtain power estimates, we used the switch-level circuit simulator IRSIM with RC-parameters extracted using EPOCH, a commercial Verilog-HDL synthesizer and standard cell router.

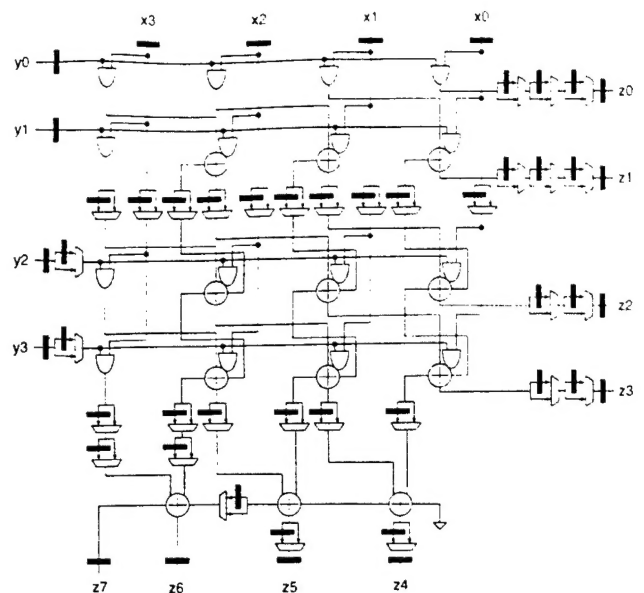


Figure 5. 4×4 array multiplier with reconfigurable 4-stage pipeline.

Figure 6 shows the relative energy savings per operation for 4-, 8-, and 16-stage reconfigurable multipliers over non-reconfigurable 4-, 8-, and 16-stage pipelined ones, respectively. For example, the second bar from the left gives the relative savings of an 8-stage reconfigurable multiplier over an 8-stage conventional multiplier when the reconfigurable multiplier is organized as a single-stage pipeline. In this case, the reconfigured 8-stage multiplier saves more than 58% of the dissipation over the conventional 8-stage multiplier. The negative savings for stages 4, 8, and 16 are due to the dissipation of the reconfiguring hardware. These numbers were obtained using IRSIM with uniformly distributed random inputs. The overall energy savings of these multipliers depend on the required maximum performance and the statistics of data rates over time.

5 Simulation Results

This section presents a comparative evaluation of a reconfigurable pipelined IQ with a conventional one on a variety of MPEG-2 MP@ML bitstreams. We first focus on one test bitstream at a fixed bit rate and provide detailed evidence about the effectiveness of our coefficient-based reconfiguration criterion. We then give simulation results that demonstrate the significant relative savings that can be achieved using our reconfigurable pipelined multiplier for a variety of bitstreams and bit rates.

The potential of reconfiguration to reduce energy dissipation can be best understood by examining the statistics




Figure 6. Relative energy reduction per operation for reconfigurable 4-, 8-, and 16-stage multipliers over non-reconfigurable 4-, 8-, and 16-stage pipelined ones, respectively.

of *flower garden*, one of the bitstreams we experimented with. Figure 7 gives the cumulative distribution of the number of nonzero quantized DCT coefficients per block for the IQ of the MPEG bitstream *flower garden* at an average bit rate of 6.0 Mbit/sec. This video is one of the test bitstreams of the MPEG video committee and consists of 38 I-pictures, 113 P-pictures, and 299 B-pictures. The resolution of each picture is 704×480 pixels. To meet the maximum throughput requirement of the IQ with low area penalty and low supply voltage, we used a reconfigurable 8-stage pipelined 16×16 multiplier operating at 31.104MHz with a 1.40V supply. Each reconfiguration mode of this multiplier is guaranteed not to change for at least 128 cycles. Figure 7 shows that 38% of the nonzero quantized DCT coefficients in the bitstream are found in blocks with at most 8 nonzero coefficients. To process each element of these blocks, it suffices to configure the 8-stage pipelined multiplier as a single-stage pipeline, thus eliminating all the intermediate register stages. The aggregate throughput of the single-stage multiplier decreases to 8 cycles per sample. For these blocks, the relative energy savings per operation for the 8-stage multiplier configured as single-stage pipeline over the non-reconfigurable 8-stage one are 58% as shown in Figure 6. Therefore, when *flower garden* is processed, the total relative savings are approximately 22.0% ($= 0.38 \times 58\%$). Similarly, 23% of the nonzero coefficients in *flower garden* can be found in blocks with 9–16 nonzero coefficients that can be processed by a two-stage configuration of the 8-stage multiplier. 31% of the nonzero coefficients in *flower garden* can be found in blocks with 17–32 nonzero coefficients that can be processed by a 4-stage con-




Figure 7. Cumulative distribution for the number of nonzero coefficients per block in *flower garden* at 6 Mbit/sec.

figuration of the 8-stage multiplier. Finally, about 0.8% of the nonzero coefficients are elements of blocks with more than 32 nonzero coefficients. In this case, the reconfigured 8-stage 16×16 multipliers should be used for maximum throughput. By adding the energy savings for each reconfiguration mode, we conclude that 42.84% are the total relative energy savings with a reconfigurable 8-stage multiplier for *flower garden*.




Figure 8. Cumulative distribution for the number of nonzero coefficients per block in *flower garden* at various bitrates.

Figure 8 gives statistics for the number of nonzero quantized DCT coefficients per block for different bit rates of the same source image *flower garden*. The percentage

of nonzero quantized DCT coefficients in blocks with a small number of nonzero DCT coefficients increases as bit rates decrease. For lower bit rates, therefore, the reconfigurable multipliers spend longer time in a "shallow", energy-efficient mode, and thus greater relative savings can be achieved. The picture size of *flower garden* at 1.5 Mbit/sec is 352×240 pixels, which is $1/4$ of the maximum picture size, 720×480 pixels, in MPEG2 MP@ML.




Figure 9. Relative energy savings with a reconfigurable 8-stage multiplier-based IQ for *susi*, *table tennis*, *mobile*, and *flower garden* at various bit rates.

Figure 9 shows the relative energy savings that can be achieved using the IQ with reconfigurable multipliers over their non-reconfigurable counterparts for the MPEG-2 MP@ML bitstreams *susi*, *table tennis*, *mobile*, and *flower garden*. Results are given for five different bit rates: 1.5, 4.0, 6.0, 8.0, and 12.0 Mbit/sec for each bitstream. Our results show that relative savings increase as video bitrates decrease and can exceed 58%.

6 Conclusion

In this paper, we have presented a novel methodology for designing reconfigurable multipliers for multimedia systems that adapt to variations in the input data rate. Our approach has been applied to the design of an inverse quantizer for MPEG-2 MP@ML and has resulted in energy reductions of up to 58%. Comparable energy savings have been observed in preliminary experiments with the inverse discrete cosine transform [3]. Our approach is also applicable to other signal processing computations and can be used to design energy-efficient pipelined arithmetic circuitry for applications with dynamically varying throughput rates.

Acknowledgments

This research was supported in part by the US Army Research Office under Grant No. DAAD19-99-1-0304.

References

- [1] B. G. Haskell, P. G. Howard, Y. A. LeCun, A. Puri, J. Ostermann, M. R. Civanlar, L. Rabiner, L. Bottou, and P. Haffner. Image and video coding-Emergin standards and beyond. *IEEE Transactions on Circuits and Systems for Video Technology*, 8(7):814-837, Nov. 1998.
- [2] S. Jou, C. Chen, E. Yang, and C. Su. A pipelined multiplier-accumulator using a high-speed, low-power static and dynamic full adder design. *IEEE Journal of Solid-State Circuits*, SC-32(1):114-118, Jan. 1997.
- [3] S. Kim and M. C. Papaefthymiou. Reconfigurable pipelining for low-energy multimedia. 2000. Unpublished Manuscript.
- [4] T. Lang, E. Musoll, and J. Cortadella. Individual flip-flops with gated clocks for low power datapaths. *IEEE Transactions on Circuits and Systems- II: Analog and Digital Signal Processing*, 44(6):507-516, June 1997.
- [5] G. Ma and F. J. Taylor. Multiplier policies for digital signal processing. *IEEE ASAP Magazine*, pages 6-19, Jan. 1990.
- [6] S. S. Mahant-Shetti, P. T. Balsara, and C. Lemonds. High performance low power array multiplier using temporal tiling. *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, 7(1):121-124, Mar. 1999.
- [7] T. H. Meng, B. M. Gordon, E. K. Tsern, and A. C. Hung. Portable video-on-demand in wireless communication. *Proceedings of The IEEE*, 83(4):659-680, Apr. 1995.
- [8] J. L. Mitchell, W. B. Pennebaker, C. E. Fogg, and D. J. LeGall. *MPEG Video Compression Standard*. Chapman and Hall, 1997.
- [9] S. R. Park and W. Burleson. Reconfiguration for power saving in real-time motion estimation. In *Proc. of 1998 IEEE International Conference on Acoustics, Speech and Signal Processing*, volume 4, pages 3037-3040, May 1998.
- [10] V. Stojanovic and V. G. Oklobdzija. Comparative analysis of master-slave latches and flip-flops for high-performance and low-power systems. *IEEE Journal of Solid-State Circuits*, SC-34(4):536-548, Apr. 1999.
- [11] N. H. E. Weste and K. E. Shraghian. *Principles of CMOS VLSI Design: A Systems Perspective*. Addison-Wesley Publishing Company, 1993.
- [12] A. Wu, C. K. Ng, and K. C. Tang. Modified booth pipelined multiplication. *Electronics Letters*, 34(12):1179-1180, June 1998.
- [13] W. Ye and M. J. Irwin. Power analysis of gated pipeline registers. In *IEEE International ASIC/SOC Conference*, pages 281-285, 1999.